

- 7 -

Bd

A capacity request is rejected by a core module if sufficient matching slots cannot be found. In order to reduce the incidence of mismatch, the matching process should always start from a selected space switch at a selected time slot and follow the same search path for each capacity request. For example, the matching process may start from space switch 0 at time slot 0 and then proceed by increasing the time slot,  $s$ , from 0 to  $T$ , where  $T$  is the number of time slots per timeframe. It then continues to the next time-port plane 53 until the 16 planes (in this example) are exhausted or the capacity is successfully allocated, whichever takes place first. The result produced by this packing search, which is well known in the art, is an occupancy pattern shown in FIG. 8.

**IN THE CLAIMS:**

Please cancel claims 1-14, and insert new claims 15-102 as follows:

- BS  
CA.
15. (New) A packet switch comprising
- a plurality of independently-controlled core modules;
  - a plurality of ingress modules each receiving packets from subtending packet sources and having a link directed to each of said core modules;
  - and
  - a plurality of egress modules each having a link from each of said core modules and transmitting packets to subtending packet sinks;
- wherein each of said ingress modules is operable to:
- issue packet-transfer requests each specifying an egress module; and
  - distribute the packet-transfer requests among said plurality of core modules for scheduling so that each of said core controllers receives a portion of said transfer requests.
16. (New) The packet switch as claimed in claim 15 further including a plurality of core controllers operating concurrently and

- 8 -

independently, one core controller associated with each of said independently-controlled core modules and having a packet scheduler.

17. (New) The packet switch as claimed in claim 15 wherein each of said ingress modules is further operable to segment each packet into data segments, each data segment occupying a time slot of a predefined duration.

18. (New) The packet switch as claimed in claim 15 wherein each ingress module has a plurality of ingress ports and an ingress controller, each of the ingress ports having an associated ingress buffer, said ingress controller being adapted to sort packets arriving at the ingress buffers from the subtending packet sources into ingress queues, each ingress queue corresponding to a one of said egress modules.

19. (New) The packet switch as claimed in claim 18 wherein said sorting is a logical sorting.

20. (New) The packet switch as claimed in claim 18 wherein the ingress controller periodically determines ingress-queue occupancy for each respective egress module and sends a capacity-request vector to a selected one of the core controllers, the capacity-request vector including capacity-allocation requests for connections to a subset of the egress modules.

21. (New) The packet switch as claimed in claim 20 wherein each of said core modules computes a schedule in response to receiving a capacity-request vector, the schedule specifying, for each capacity request, time slots in a predefined time frame.

22. (New) The switch as claimed in claim 21 wherein each core module comprises a single-stage space switch.

23. (New) The switch as claimed in claim 22 wherein said space switch is a rotator space switch comprising an input rotator, a plurality of memory devices, and an output rotator.

- 9 -

24. (New) The switch as claimed in claim 21 wherein the core modules, the ingress modules, and the egress modules are spatially distributed.

25. (New) The switch as claimed in claim 24 wherein each core module comprises a number  $S \geq 1$  of single-stage space switches, each space switch having a plurality of input ports each connecting to an ingress module and a plurality of output ports each connecting to an egress module.

26. (New) The switch as claimed in claim 25 wherein said single-stage space switch is an optical switch.

27. (New) The switch as claimed in claim 25 wherein each ingress module has a plurality of timing circuits each communicating with a time counter associated with a one of said core modules to realize time coordination between said each ingress module and said one of said core modules.

28. (New) The packet switch as claimed in claim 25 wherein each link carries a wavelength-division-multiplexed optical signal comprising  $S$  wavelength channels and wherein each of said single-stage space switches channels of the same wavelength.

29. (New) A method of scheduling performed by a controller of a core module having at least one link from each of a plurality of ingress modules, said core module comprising a plurality of  $S \geq 1$  space switches, each space switch having a plurality of input ports and a plurality of output ports, the method comprising steps of:

creating a data structure comprising:

a first three-dimensional matrix having a space dimension  $s$  representing space switches associated with the core module, a space dimension  $p$  representing space-switch input ports, and a time dimension  $t$  representing time slots in a slotted time frame; and

- 10 -

a second three-dimensional matrix having said space dimension  $s$ , a space dimension  $\pi$  representing space-switch output ports, and said time dimension  $t$ ;

receiving capacity-allocation requests from the ingress modules, each request specifying an input port  $p$ , an output port  $\pi$  and a number  $K$  of time slots per time frame; and

selecting a space switch  $s$  and a time slot  $t$  and, if both entries  $\{s, p, t\}$  of the first matrix and  $\{s, \pi, t\}$  of the second matrix are free, allocating space-switch  $s$  and time-slot  $t$  and marking entries  $\{s, p, t\}$  and  $\{s, \pi, t\}$  as busy.

30. (New) The method as claimed in claim 29 further including a step of repeating said selecting until at most  $K$  time slots are allocated.

31. (New) The method as claimed in claim 29 further including a step of terminating a current connection by setting the value of  $K$  to zero.

32. (New) The method as claimed in claim 30 wherein said selecting for each capacity-allocation request considers all time slots in said time frame, then all of said plurality of space switches.

33. (New) The method as claimed in claim 32 further including a step of producing a scheduling matrix that associates each of said input ports with each of said output ports during each of said time slots.

34. (New) The method as claimed in claim 33 wherein said all time slots are considered in a predetermined order.

35. (New) The method as claimed in claim 33 wherein said all of said plurality of switches are considered in a predetermined order.

36. (New) The method as claimed in claim 33 including the further steps of:

- 11 -

assembling said capacity-allocation requests into a capacity-request matrix each entry of which containing a requested capacity allocation; and

attempting to schedule each entry in the matrix to perform core reconfiguration.

37. (New) The method as claimed in claim 36 including the further step of maintaining at each ingress module two scheduling matrices, one in current use and one in update mode.

38. (New) The method as claimed in claim 37 wherein each time a core reconfiguration occurs, a scheduling matrix in use is swapped for a current scheduling matrix.

39. (New) The method as claimed in claim 38 wherein an unused copy of the scheduling matrix is available for update after the core reconfiguration.

40. (New) The method as claimed in claim 39 wherein rows in the scheduling matrix are executed sequentially, one per time slot, until a next core module reconfiguration occurs, and, after core module reconfiguration, processing continues at a next time slot.

41. (New) A distributed packet switch comprising:

a plurality of  $m$  cross connectors each having  $n$  outer links and  $n$  inner links;

a plurality of  $n$  core modules each core module comprising a number of space switches;

a plurality of  $m \times n$  edge modules; and

a plurality of  $n$  core controllers, one core controller associated with each of said core modules;

wherein

- 12 -

each of said outer links connects to an edge module and includes  $\Lambda$  channels in each direction to and from the edge module;

each of said inner links connects to a core module and includes  $\Lambda$  channels in each direction to and from the core module; and

said number of space switches per core module does not exceed the ratio  $\Lambda/n$ .

42. (New) The distributed packet switch as claimed in claim 41 further including an edge controller associated with each of said edge modules;

43. (New) The distributed packet switch as claimed in claim 42 wherein said edge modules and said core modules are spatially distributed over a wide geographical area.

44. (New) The distributed packet switch as claimed in claim 43 wherein each of said outer links is a wavelength-division-multiplexed link and each of said inner links is a wavelength-division-multiplexed link.

45. (New) The distributed packet switch as claimed in claim 43 wherein each of said edge modules has means for time coordination with each of said core modules.

46. (New) The distributed packet switch as claimed in claim 45 wherein said means includes a timing circuit.

47. (New) The distributed packet switch as claimed in claim 46 wherein each of said core modules has means for time coordination with each of said edge modules.

48. (New) The distributed packet switch as claimed in claim 47 wherein said means includes a timing circuit.

- 13 -

49. (New) The distributed packet switch as claimed in claim 48 wherein at least one of said space switches in each of said core modules operates in a time-division-multiplexed mode.

50. (New) The distributed packet switch as claimed in claim 49 wherein at least two edge modules transmit time-division-multiplexed signals to said at least one of said space switches and adjust their transmission times so that the time-division-multiplexed signals are received at said at least one of said space switches in time alignment.

51. (New) The distributed packet switch as claimed in claim 50 wherein one of said edge modules is collocated with a selected one of said core modules and hosts a controller that serves as the core controller of said one of said core modules.

52. (New) The distributed packet switch as claimed in claim 50 wherein each of said edge modules is adapted to transmit capacity-allocation requests to any of said core modules.

53. (New) The distributed packet switch as claimed in claim 52 wherein the core controller associated with said any of said core modules computes a schedule in response to receiving said capacity-allocation requests, the schedule specifying, for each capacity request, time slots in a predefined time frame.

54. (New) The distributed packet switch as claimed in claim 53 wherein at least one of said space switches is an electronic space switch.

55. (New) The distributed packet switch as claimed in claim 53 wherein at least one of said space switches is an optical space switch.

56. (New) The distributed packet switch as claimed in claim 53 wherein at least one of said space switches is an electronic single-stage rotator switch.

57. (New) A packet switch comprising:

BS  
Cost

- 14 -

a plurality of egress modules, each for transmitting packets on at least one network link;

a plurality of ingress modules, each for receiving packets from at least one network link and capable of making ingress-to-egress-module connection requests for transferring received packets to any other of the egress modules; and

a plurality of core modules, each capable of simultaneously receiving and independently responding to said connection requests from any of the ingress modules, and of providing ingress-to-egress-module connections between any of the ingress modules and any of the egress modules in response to said connection requests.

35  
CotA.  
58. (New) The packet switch as claimed in claim 57 wherein each core module has its own controller for allocating and scheduling resources to the ingress-to-egress-module connections, said own controller operating independently of the other core modules' controllers.

59. (New) The packet switch as claimed in claim 57 wherein each ingress module is connected to each core module by at least one link and each core module is connected to each egress module by at least one link.

60. (New) The packet switch as claimed in claim 59 wherein said resources comprise data paths of sufficient capacities to accommodate said connection requests.

61. (New) The packet switch as claimed in claim 60 wherein said controller of said each core module operates concurrently with the controllers of the other core modules.

62. (New) The packet switch as claimed in claim 61 wherein each core module includes circuitry for time domain multiplexing a plurality of ingress-to-egress-module connections.



- 15 -

63. (New) The packet switch as claimed in claim 62 wherein each of the plurality of ingress-to-egress-module connections starts at the same ingress module and ends at the same egress module.

64. (New) The packet switch as claimed in claim 63 wherein the ingress modules include means for routing a connection through more than one of said core modules.

65. (New) The packet switch as claimed in claim 64 additionally including at least one cross-connector connecting a subset of the ingress modules to the core modules.

66. (New) The packet switch as claimed in claim 65 wherein all of said core modules switch data blocks of the same time duration.

67. (New) The packet switch as claimed in claim 66 wherein each of said core modules maintains its own time reference.

68. (New) The packet switch as claimed in claim 67 wherein each of said ingress modules times the transmission of its data blocks to arrive at a selected one of said core modules at a time determined by said selected one of said core modules.

69. (New) The packet switch as claimed in claim 57 wherein at least one of said ingress modules is provided with a table of preferred paths to each of said egress modules, each path being defined by a core module.

70. (New) The packet switch as claimed in claim 66 wherein at least one of said core modules comprises at least one space switch.

71. (New) The packet switch as claimed in claim 66 wherein at least one of said core modules is a single-stage electronic rotator switch.

72. (New) The packet switch as claimed in claim 66 wherein each ingress module has a plurality of ingress ports and an ingress controller, each of said ingress ports including:

- 16 -

an associated ingress buffer for receiving packets from subtending packet sources; and

means for sorting packets arriving in the ingress buffer into ingress queues each ingress queue corresponding to an egress module from which packets in said each queue are to egress from the switch for delivery to the subtending packet sinks.

73. (New) The packet switch as claimed in claim 72 wherein each ingress module is further adapted to send capacity-allocation requests to a one of said core modules.

74. (New) The packet switch as claimed in claim 73 wherein each ingress module is further adapted to receive connection schedules from each of said core modules.

75. (New) The packet switch as claimed in claim 73 wherein each of said connections is scheduled for transmission over time-slots of equal duration and each of said connection schedules includes time-slot identifiers for each scheduled connection.

76. (New) The packet switch as claimed in claim 73 wherein each ingress module is further adapted to:

create a vector of pointers to the sorted packets; and

assemble said connection schedules into a scheduling matrix so that a non-blank entry in the scheduling matrix indicates an index of the vector of pointers.

77. (New) The packet switch as claimed in claim 76 wherein said ingress controller receives capacity requirements from subtending sources and determines said capacity-allocation requests.

78. (New) The packet switch as claimed in claim 76 wherein the ingress controller periodically determines a number of packets waiting in the

- 17 -

ingress queues for each respective egress module and determines said capacity-allocation requests.

79. (New) The packet switch as claimed in claim 75 wherein each core module comprises a plurality of single-stage rotator switches, each rotator switch having a number of input ports collectively adapted to accommodate a number of channels at least equal to the number of ingress modules and a number of output ports collectively adapted to accommodate a number of channels at least equal to the number of egress modules, each ingress module having at least one channel to each of the rotator switches, and each egress module having at least one channel from each of the rotator switches.

80. (New) The packet switch as claimed in claim 79 wherein each ingress module is combined with an egress module to form an integrated edge module.

81. (New) The packet switch as claimed in claim 75 wherein the core modules are co-located at one geographical location.

82. (New) The packet switch as claimed in claim 75 wherein the core modules, the ingress modules, and the egress modules are spatially distributed.

83. (New) The packet switch as claimed in claim 75 wherein each ingress module has a number of timing circuits at least equal to the number of core modules, each of the timing circuits being time-coordinated with a time counter associated with each of said core modules.

84. (New) The packet switch as claimed in claim 83 wherein one edge module is co-located with each core module and said one edge module serves as a controller for the core module.

85. (New) A method of switching a packet through a switch comprising a plurality of ingress modules each having at least one ingress port, a plurality of egress modules each having at least one egress port, and a

- 18 -

plurality of core modules, wherein each ingress module is coupled to each core module, each core module is coupled to each egress module, and packets can traverse only one ingress module, one core module and one egress module in moving from an ingress port to an egress port, the method comprising steps of:

receiving a packet from a subtending traffic source at one of the ingress modules;

selecting at the ingress module one of the egress modules to which to send the packet;

sending from the ingress module a connection request to a selected one of the core modules, requesting a connection of a specified capacity to said one of the egress modules; and

at said selected one of the core modules:

determining a feasible capacity allocation,

subtracting said feasible capacity allocation from said specified capacity to update said connection request, and

returning the connection request to said one of the ingress modules.

86. (New) The method as claimed in claim 85 including the further step of said one of the ingress modules sending the connection request to another of the core modules if said feasible capacity allocation is less than said specified capacity.

87. (New) The method as claimed in claim 86 wherein said any one of the core modules is selected at random.

88. (New) The method as claimed in claim 86 wherein said any one of the core modules is selected according to a preferred order.

- 19 -

89. (New) The method as claimed in claim 88 wherein said preferred order is specific to said one of the ingress modules and said one of the egress modules.

~~90.~~ (New) An ingress module in a packet switch, the ingress module comprising:

an ingress controller;

a plurality of ingress ports each having an ingress buffer for receiving packets from subtending packet sources, each packet indicating a one of predefined destinations;

a plurality of output ports for directing said packets to a plurality of core modules;

means for sorting the packets received in said ingress buffer into ingress queues each corresponding to one of said destinations;

means for storing a set of predefined paths to each of said predefined destinations;

means for formulating connection requests, each of said connection requests specifying a destination and a required capacity allocation; and

means for selecting a candidate path from among said predefined paths for each connection request.

91. (New) The ingress module as claimed in claim 90 wherein said ingress controller performs said sorting, storing, formulating, and selecting.

92. (New) The ingress module as claimed in claim 91 further including means for distributing said connection requests among said predefined paths.

93. (New) The ingress module as claimed in claim 92 further including means for selecting more than one path for a connection.

- 20 -

94. (New) The ingress module as claimed in claim 91 wherein each of said predefined paths is defined by one of said core modules.

95. (New) The ingress module as claimed in claim 94 wherein said core modules are distributed over a wide area.

96. (New) The ingress module as claimed in claim 95 further including a timing circuit adapted to perform time coordination with each of said core modules.

97. (New) A core module in a packet switch, the core module comprising:

at least one space switch having a plurality of input ports and a plurality of output ports;

a core controller adapted to receive connection requests, each connection request specifying a required capacity allocation and a destination selected from among a set of predefined destinations;

means for associating each destination with one of said output ports;

a scheduler associated with said core controller, said scheduler adapted to time the transfer of packets from said input ports to said output ports; and

means for communicating with sources of said connection requests.

98. (New) The core module as claimed in claim 97 wherein said space switch is a bufferless switch operating in a time-multiplexed mode.

99. (New) The core module as claimed in claim 98 wherein said sources are ingress modules each provided with timing circuitry to control the timing of packet transfer.

100. (New) The core module as claimed in claim 99 further including a timing circuit adapted for time coordination with said timing circuitry.

- 21 -

BS  
Concl.

101. (New) The core module as claimed in claim 100 further adapted to receive wavelength-division-multiplexed channels and assign a wavelength channel to each of said at least one space switch.

102. (New) The core module as claimed in claim 101 wherein said core controller determines a feasible capacity allocation for a connection request not exceeding said required capacity allocation and, when the feasible capacity allocation is less than the required capacity allocation, modifies the connection request and returns a modified connection request to a source of the connection request.

---